

Symbiosis of Search and Heuristics for Random 3-SAT

Sid Mijnders, Boris de Wilde, and Marijn J. H. Heule*

Delft University of Technology, Delft, The Netherlands

Abstract. When combined properly, search techniques can reveal the full potential of sophisticated branching heuristics. We demonstrate this observation on the well-known class of random 3-SAT formulae. First, a new branching heuristic is presented, which generalizes existing work on this class. Much smaller search trees can be constructed by using this heuristic. Second, we introduce a variant of discrepancy search, called ALDS. Theoretical and practical evidence support that ALDS traverses the search tree in a near-optimal order when combined with the new heuristic. Both techniques, search and heuristic, have been implemented in the look-ahead solver *march*. The SAT 2009 competition results show that *march* is by far the strongest complete solver on random k -SAT formulae.

1 Introduction

Satisfiability (SAT) solvers have a rich history of branching heuristics [1]. These heuristics are crucial for fast performance. They can be split into *decision heuristics* and *direction heuristics*. The former selects for each node in the tree a *decision variable* to branch on. These decision heuristics determine the size of the search tree. For each free value of the decision variable a child node is created. Direction heuristics provide the preferred order in which these child nodes should be visited. In case a search tree contains solutions, effective direction heuristics can boost performance.

Two techniques are designed to repair errors made by branching heuristics. First, *restart strategies* [2] are used to compensate for ineffective choices made by decision heuristics. A new search tree is created after every restart. Second, *discrepancy search* [3], a technique used in Constraint Programming, focuses on those parts of the search tree that mostly follow the preference of the direction heuristics. This technique only changes the order in which the search tree is traversed.

In this paper, we present a new branching heuristic for look-ahead SAT solvers, called the *recursive weight heuristic*. This heuristic is a generalization of earlier look-ahead evaluation heuristics [4–6]. Regarding the decision heuristic, we show that on random 3-SAT formulae the search tree is significantly smaller compared to alternative heuristics. Although the heuristic is quite expensive in terms of computational costs, performance is clearly improved.

Also, the new direction heuristic results in an observable bias in the distribution of solutions. To capitalize on this, we developed a new discrepancy search algorithm, called advanced limited discrepancy search (ALDS), which combines features of improved limited discrepancy search [7] and depth-bounded discrepancy search [8]. We

* Supported by the Dutch Organization for Scientific Research (NWO) under grant 617.023.611

provide both theoretical and experimental evidence to show that the combination of ALDS and the recursive weight heuristic, traverse the tree in a near-optimal order on random 3-SAT formulae.

The outline of the paper is as follows: First, we will explain look-ahead heuristics in Section 2, both the existing work and our new heuristic. In Section 3, various search techniques will be discussed. The focus will be discrepancy search and our variant of this technique. Section 4 will offer theoretical and practical results showing that the combination of ALDS and our heuristic is effective on random 3-SAT instances. Finally, we draw some conclusions in Section 5.

2 Look-ahead heuristics

Most work on branching heuristics in the field of Satisfiability focuses on look-ahead SAT solvers [1]. In contrast to many other solvers, look-ahead solvers keep track of various statistical measurements that make it possible to use quite complex heuristics. In this section, we first will provide an overview of look-ahead SAT solvers. Afterwards, the branching heuristics in these solvers are discussed. We conclude this section by introducing an improved heuristic.

2.1 Look-ahead SAT solvers

The look-ahead architecture for SAT solvers is based on the *DPLL* framework [9]: It is a complete solving method which selects in each step a decision variable x_{decision} and recursively calls DPLL for the reduced formula where x_{decision} is assigned to false (denoted by $\mathcal{F}[x_{\text{decision}} = 0]$) and another where x_{decision} is assigned to true (denoted by $\mathcal{F}[x_{\text{decision}} = 1]$).

A formula \mathcal{F} is reduced by *unit propagation*: Given a formula \mathcal{F} , an unassigned variable x and a Boolean value \mathbf{B} , first x is assigned to \mathbf{B} . If this assignment φ results in a *unit clause* (clause of size 1) then φ is expanded by assigning the remaining literal of that clause to true. This is repeated until no unit clauses are left in φ applied to \mathcal{F} . We denote by $\varphi * \mathcal{F}$ the reduced formula after unit propagation of applying φ on \mathcal{F} , with all satisfied clauses removed. So, more specific than above, $\mathcal{F}[x = \mathbf{B}] := \varphi * \mathcal{F}$.

Algorithm 1 UNITPROPAGATION(formula \mathcal{F} , variable x , $\mathbf{B} \in \{0, 1\}$)

```

1:  $\varphi := \{x \leftarrow \mathbf{B}\}$ 
2: while empty clause  $\notin \varphi$  applied on  $\mathcal{F}$  and unit clause  $(y) \in \varphi$  applied on  $\mathcal{F}$  do
3:    $\varphi := \varphi \cup \{y \leftarrow 1\}$ 
4: end while
5: return  $\varphi * \mathcal{F}$ 

```

The recursion has two kinds of leaf nodes: Either all clauses have been satisfied (denoted by $\varphi * \mathcal{F} = \emptyset$), meaning that a satisfying assignment has been found, or $\varphi * \mathcal{F}$ contains an *empty clause* (a clause of which all literals have been falsified), meaning a dead end. In the latter case the algorithm backtracks.

The core of the look-ahead architecture is the LOOKAHEAD procedure, which incorporates the branching heuristics (selecting a decision variable and selecting the first branch) and several reasoning techniques to reduce the size of the formula. Because the latter are beyond the scope of this paper, we refer the reader to [10] for details. Algorithm 2 shows the top level structure. Notice that the LOOKAHEAD procedure returns a reduced formula \mathcal{F} , variable x_{decision} , and value \mathbf{B} . Fig. 1 provides a graphical overview of the architecture.

Algorithm 2 DPLL(formula \mathcal{F})

```

1: if  $\mathcal{F} = \emptyset$  then
2:   return satisfiable
3: end if
4:  $\langle \mathcal{F}; x_{\text{decision}}; \mathbf{B} \rangle := \text{LOOKAHEAD}(\mathcal{F})$ 
5: if empty clause  $\in \mathcal{F}$  then
6:   return unsatisfiable
7: end if
8: if DPLL(  $\mathcal{F}[x_{\text{decision}} = \mathbf{B}]$  ) = satisfiable then
9:   return satisfiable
10: end if
11: return DPLL(  $\mathcal{F}[x_{\text{decision}} = \neg \mathbf{B}]$  )

```

The LOOKAHEAD procedure, as the name suggests, performs *look-aheads*. A look-ahead on x starts by assigning x to true followed by unit propagation. The importance of x is measured *and* possible reductions of the formula are detected. After this analysis, it backtracks, ending the look-ahead. The rationale of a look-ahead operation is that evaluating the effect of actually assigning variables to truth values and performing unit propagation is more adequate than taking a cheap guess using statistical data on \mathcal{F} .

2.2 Look-ahead evaluation

Branching heuristics in look-ahead SAT solvers are based on evaluating the reduction of the formula during a look-ahead. This reduction is expressed using the *difference* or *distance heuristic* (in short DIFF). The larger the reduction, the higher the heuristic value. A DIFF could be based on many statistics, such as the reduction of the number of variables. Yet, all look-ahead SAT solvers use a DIFF based only on the set of newly created (i.e. reduced, but not satisfied) clauses, denoted by C_{new} [10].

The decision variable x_{decision} is selected by combining for each variable x_i the values $\text{DIFF}(\mathcal{F}, \mathcal{F}[x_i = 0])$ and $\text{DIFF}(\mathcal{F}, \mathcal{F}[x_i = 1])$. The objective of the decision heuristic is to construct a small and balanced search tree. The product of these numbers is generally considered to be an effective heuristic for this purpose [1]. The sum of these numbers can be used for tie-breaking.

Once x_{decision} is selected, the direction heuristics decide whether to branch first on $\mathcal{F}[x_{\text{decision}} = 0]$ or $\mathcal{F}[x_{\text{decision}} = 1]$. Most solvers prefer the branch which is the most satisfiable [1, 10]. A heuristic used to determine the most satisfiable branch selects Boolean value \mathbf{B} for which $\text{DIFF}(\mathcal{F}, \mathcal{F}[x_i = \mathbf{B}])$ is the smallest.

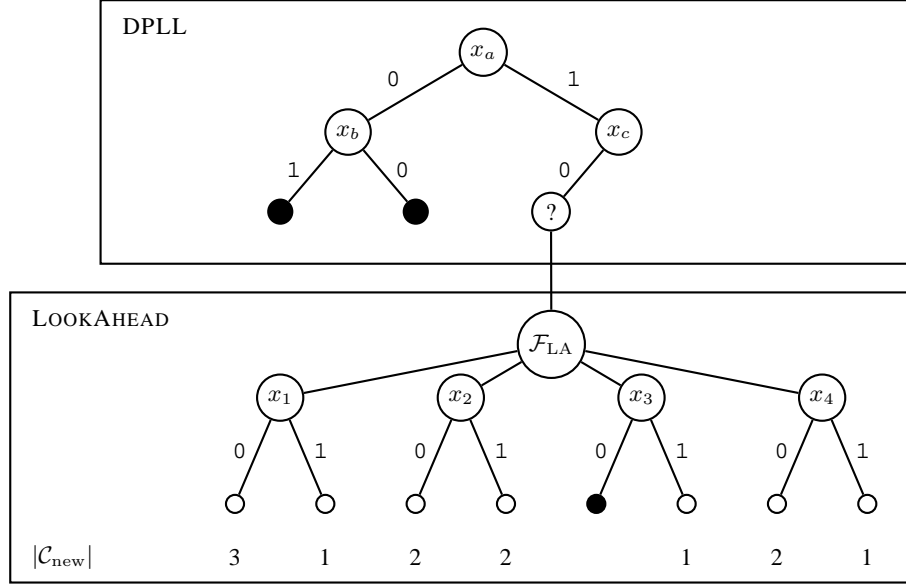


Fig. 1. A graphical representation of the look-ahead architecture. Above, the DPLL superstructure (a binary tree) is shown. In each node of the DPLL-tree, the LOOKAHEAD procedure is called to select the decision variable and to compute implied variables by additional reasoning. Black nodes refer to leaf nodes.

Consider the following example formula:

$$\mathcal{F}_{LA} = (\neg x_1 \vee x_3) \wedge (x_1 \vee x_2 \vee x_3) \wedge (x_1 \vee \neg x_2 \vee x_4) \wedge (x_1 \vee \neg x_2 \vee \neg x_4) \wedge (x_2 \vee \neg x_3 \vee x_4)$$

Since all clauses in \mathcal{F}_{LA} have size three or smaller, only new binary clauses can be created. For instance, during the look-ahead on $\neg x_1$, three new binary clauses are created (all clauses in which literal x_1 occurs). The look-ahead on x_1 will force x_3 to be assigned to true by unit propagation. This will reduce the last clause to a binary clause, while all other clauses become satisfied. Similarly, we can compute the number of new binary clauses for all look-aheads – see Fig. 1.

Finally, the selection of the decision variable is based on the reduction measurements of both the look-ahead on $\neg x_i$ and x_i . Generally, the product is used to combine the numbers. In this example, x_2 would be selected as decision variable, because the product of the reduction measured while performing look-ahead on $\neg x_2$ and x_2 is the highest (i.e. 4).

The DIFF heuristic based on the number of newly created clauses was introduced by Li and Anbulagan [4]. Several extensions have been proposed dealing with how to weigh clauses in \mathcal{C}_{new} . If the original formula is 3-SAT, \mathcal{C}_{new} only consists of binary clauses. For this special case, Li [5] uses in **satz** weights based on the occurrences of variables. Let $\#(x)$ denote the number of occurrences of literal x . Each clause $(y \vee z) \in$

\mathcal{C}_{new} gets a weight of $\#(\neg y) + \#(\neg z)$. A slight variation but much more effective weight is used by Dubois and Dequen [6] in their solver *kcnfs*. Their backbone search heuristic weighs a binary clause $(y \vee z)$ by $\#(\neg y) \times \#(\neg z)$. In case of k -SAT instances, Kullmann [11] uses in the *OKsolver* weights based on the length of clauses in \mathcal{C}_{new} . A clause of size k roughly get a weight of $\gamma_k = 5^{2-k}$.

2.3 Recursive weight heuristic

We developed a model to generalize existing work on look-ahead branching heuristics. We refer to this model as the *recursive weight heuristic*. Let $\text{VAR}(\mathcal{F})$ refer to the set of variables in \mathcal{F} and $n = |\text{VAR}(\mathcal{F})|$. The heuristic values $h_i(x)$ express for each iteration i how much literal x is forced to true by the clauses containing x . First, for all literals x , $h_0(x)$ are initialized on 1:

$$h_0(x) = h_0(\neg x) = 1 \quad (1)$$

At each step, the heuristics values $h_i(x)$ are scaled using the average value μ_i :

$$\mu_i = \frac{1}{2n} \sum_{x \in \text{VAR}(\mathcal{F})} (h_i(x) + h_i(\neg x)) \quad (2)$$

Finally, in each next iteration, the heuristic values $h_{i+1}(x)$ are computed in which literals y get weight $h_i(\neg y)/\mu_i$. Weight γ expresses the relative importance of binary clauses. This weight could also be seen as the heuristic value of a falsified literal.

$$h_{i+1}(x) = \sum_{(x \vee y \vee z) \in \mathcal{F}} \left(\frac{h_i(\neg y)}{\mu_i} \cdot \frac{h_i(\neg z)}{\mu_i} \right) + \gamma \sum_{(x \vee y) \in \mathcal{F}} \frac{h_i(\neg y)}{\mu_i} \quad (3)$$

Earlier work on look-ahead heuristics can be formulated using the model above. For binary clause $(y \vee z) \in \mathcal{C}_{\text{new}}$ they compute a weight $w_i^*(y \vee z) = h_i(\neg y) * h_i(\neg z)$:

- Li & Anbulagan 1997 [4]: $w(y \vee z) = 1 = h_0(\neg y) \times h_0(\neg z)$, in short w_0^\times .
- Li 1999 [5]: $w(y \vee z) = \#(\neg y) + \#(\neg z) = h_1(\neg y) + h_1(\neg z)$, in short w_1^+ .
- Dubois 2001 [6]: $w(y \vee z) = \#(\neg y) \times \#(\neg z) = h_1(\neg y) \times h_1(\neg z)$, in short w_1^\times .

Although [5,6] used $\gamma = 5$, we observed stronger performance using $\gamma = 3.3$. Also, the size of the tree can be reduced significantly by using weights w_i^\times with $i > 1$.

We implemented $w_0^\times, w_1^+, w_1^\times, w_2^\times, w_3^\times$ and w_4^\times in the look-ahead solver *march_ks* [12] with $\gamma = 3.3$. We experimented on 500 random 3-SAT formulae with 450 variables and 1915 clauses (phase transition density). To provide stable numbers, all instances were unsatisfiable. Fig. 2 shows the results. Clearly, the average size of the tree is smaller using w_3^\times compared to the alternative heuristics. Although w_3^\times is much more expensive to compute, the average time to solve these instances has also decreased. Regarding the computational costs, we observed that w_3^\times resulted in best performance of all w_i^\times on random 3-SAT formulae. In case $i > 3$, the reduction of the size of the tree is not large enough to compensate for the additional cost to compute the weights.




Fig. 2. The influence of different heuristics for the clause weight used for look-ahead evaluation on the average size of the search tree and the computational costs. The results are average sizes and times for 500 random 3-SAT formulae with 450 variables and 1915 clauses.

Our implementation with w_3^x participated as **march_hi** at the SAT competition of 2009¹. It won the random unsatisfiable category. Apart from a few minor optimizations and fixes, the only difference compared to **march_ks** (the winner in 2007) is the recursive weight heuristic. Both versions competed during SAT 2009 and **march_hi** solved over 10% more unsatisfiable instances. For many instances in this category it was the only program to solve them.

3 Heuristic search

Let us take a step back from SAT to consider how to capitalize on direction heuristics in general. We will first discuss the terminology of direction heuristics, and the ideas behind the most important heuristic search strategies. Afterwards, we will introduce an alternative search strategy, which is very powerful in combination with the recursive weight heuristic.

3.1 Direction heuristics

There are various ways to explore search trees. Searching the entire tree for a specific *goal node* is costly. Therefore, *search strategies* have been developed to guide the search towards a goal node. To show that a problem has no solutions, the search has to be *complete* by visiting all leaf nodes. Complete search strategies will either detect a goal node or prove that none exists.

If a search tree contains goal nodes, *direction heuristics* predict which branches have a higher probability of leading to a goal node than others. The branch with the highest preference will be called the *left branch*. Any other branch is called a *discrepancy*. In the case of a binary search tree a discrepancy can also be referred to as the *right branch*.

¹ see <http://www.satcompetition.org> for details

In theory, direction heuristics are very powerful. Perfect direction heuristics would lead to the goal node immediately. If such a perfect direction heuristic would exist, given that it is computable in polynomial time, it would prove that $\mathcal{P} = \mathcal{NP}$. At each node direction heuristics have a probability of picking the correct branch as left branch. This probability is called the *heuristic probability*. Another probability that we define is the *goal node probability*. For each node it expresses the probability that the subtree with this node as root contains a goal node. The heuristic probability $P_{\text{heur}}(v)$ and the goal node probability $P_{\text{goal}}(v)$ are related as follows:

$$\begin{aligned} P_{\text{goal}}(v) &= P(\text{goal-node}(v)) \\ P_{\text{heur}}(v) &= P(\text{goal-node}(\text{left-child}(v)) \mid \text{goal-node}(v)) \\ P_{\text{goal}}(\text{left-child}(v)) &= P(\text{goal-node}(\text{left-child}(v))) \\ &= P_{\text{goal}}(v) \cdot P_{\text{heur}}(v) \end{aligned}$$

In the case of the binary search tree we can also define:

$$\begin{aligned} P_{\text{goal}}(\text{right-child}(v)) &= P(\text{goal-node}(\text{right-child}(v))) \\ &= P_{\text{goal}}(v) \cdot (1 - P_{\text{heur}}(v)) \end{aligned}$$

For a given dataset D and search strategy S , $P_{\text{goal}}(v)$ denotes the fraction of formulae in D that contain solutions in the subtree rooted at v , while applying algorithm S .

It is observed [12], that heuristics tend to make more mistakes in the top of the search tree. As we get closer to the leaf nodes in the tree, the underlying problem has been simplified. Heuristics perform better on simplified problems. Therefore, it is expected that $P_{\text{heur}}(v)$ increases while descending in the search tree. Consequently, direction heuristics are most likely to make mistakes near the root of the search tree.

To illustrate goal node probabilities throughout a tree, we will consider a direction heuristic with increasing $P_{\text{heur}}(v)$ in a binary tree with a single solution, see Fig. 3. Notice that, if a problem has solutions, $P_{\text{goal}}(\text{root}) = 1$.

Looking at the goal node probabilities in this example, it shows that, although for each node individually the left child has a higher $P_{\text{goal}}(v)$ than the right child, when comparing all nodes at a certain depth no clear pattern can be observed between the $P_{\text{goal}}(v)$ values. While searching for the goal node, one wants to take the $P_{\text{goal}}(v)$ of the leaf nodes into account. In this example, when a search strategy visits the nodes with a high goal node probability first, it will on average visit less leaf nodes before finding a goal node. We will now discuss several complete search strategies.

3.2 Depth first search

One of the best known search strategies is *depth first search* (DFS). DFS branches left until it reaches a leaf node, after which it backtracks chronologically. The order in which DFS visits leaf nodes, from left to right, is shown in Fig. 4. DFS traverses the minimum number of edges needed to explore the entire tree. When $P_{\text{heur}}(v) = 0.5$ for all nodes v in a binary tree, meaning that the direction heuristic might as well randomly select branches, $P_{\text{goal}}(v)$ is equal for every node v at the same depth, then DFS is the cheapest strategy to use. When direction heuristics are stronger than random selection, as in Fig. 3, alternative strategies traverse the tree more efficiently.




Fig. 3. Binary search tree containing one goal node. Nodes are labelled with the goal node probability $P_{\text{goal}}(v)$. Edges show the heuristic probabilities $P_{\text{heur}}(v)$ (left) and $1 - P_{\text{heur}}(v)$ (right).




Fig. 4. The iterations of DFS on a search tree of depth 3.

3.3 Discrepancy search

Assume that direction heuristics perform significantly better than random branch selection. When some leaf node is not a goal node, the intuition is that the heuristic only took a small number of wrong branches on the path to this leaf node. Intuitively, tree search performs best when paths with a small number of discrepancies are explored first.

Limited Discrepancy Search (LDS) LDS [3] explores first those parts of the search tree that have a small number of discrepancies. In each iteration of LDS the number of allowed discrepancies is incremented. Fig. 5 illustrates the iterations of LDS. LDS has some redundancy since it only sets an upper bound on the allowed number of discrepancies. Therefore, in iteration x , LDS examines the paths from previous iterations again.

Improved Limited Discrepancy Search (ILDS) ILDS improves LDS by eliminating the redundancy. This is achieved by providing a maximum search depth d to the algorithm. Given this depth, at any point during its execution, the algorithm keeps track of

the remaining number of depths to be searched. As a consequence, in each iteration x , only the paths with exactly x discrepancies are explored (starting with $x = 0$). This way ILDS ensures that subtrees rooted at depth d are explored only once. All subtrees rooted at depth d are searched using DFS. The iterations of ILDS are shown in Fig. 6.




Fig. 5. The iterations of LDS on a search tree of depth 3.




Fig. 6. The iterations of ILDS on a search tree of depth 3.

Depth-bounded Discrepancy Search (DDS) By incrementally increasing the maximum depth up to which discrepancies are allowed to occur, DDS [8] differs from (I)LDS. More specific, DDS visits in each iteration $i + 1$, all branches at depth $d < i$, only the discrepancies at depth $d = i$, and no discrepancies are allowed for $d > i$. Exploring the search tree this way also removes the requirement of specifying a maximum depth. As can be seen in Fig. 7, DDS explores paths with multiple right branches at the top of the search tree relatively early. In specific cases where the direction heuristics are bad (heuristic probabilities are close to 0.5 in case of a binary tree) near the root of the tree, but suddenly get very good (close to 1) at a certain depth, it is useful to introduce multiple discrepancies at the top of the search tree early.




Fig. 7. The iterations of DDS on a search tree of depth 3.

3.4 Advanced Limited Discrepancy Search

The papers describing LDS [3], ILDS [7] and DDS [8] are precise on which leaf nodes are explored in each of the iteration stages. Yet, apart from the pseudocode, the order in which leaf nodes are visited within a single iteration stage is not explicitly specified. We will assume that the strategies are applied as described in the pseudocode. For ILDS, DDS this means that, in each iteration, leaf nodes are explored from left to right.

By combining features from ILDS and DDS, a new search strategy can be created. This search strategy uses the iterations of ILDS, while nodes within an iteration are visited according to DDS. More specific, nodes with the same number of discrepancies are visited from right to left. We call this strategy *Advanced Limited Discrepancy Search* (ALDS). The iterations of this strategy are shown in Fig. 8. Like ILDS, ALDS is only applied until a certain depth d , while subtrees rooted at d are explored using DFS.




Fig. 8. The iterations of ALDS on a search tree of depth 3.

This search strategy is inspired by an earlier experimental study on random 3-SAT instances, where we observed [12] two patterns regarding the goal node probabilities:

1. Leaf nodes reached with less discrepancies have a higher goal node probability.
2. For leaf nodes reached with the same number of discrepancies, those reached with discrepancies closer to the root have a higher goal node probability.

Notice that ALDS visits leaf nodes in the preferred order of these observations.

In order to compare the various search strategies, we propose a model to approximate the performance. In this model, the top of the search tree, until depth d , is visited using discrepancy search, while all subtrees rooted at depth d are visited by DFS. The depth d at which DFS takes over from the discrepancy search is called the *jump depth*. So, using a jump depth d would result in 2^d subtrees. Discrepancy search ensures that promising parts of the search tree are explored first, while DFS searches the remaining subtree with minimal branching overhead. Subtrees explored by DFS are considered leaf nodes in the discrepancy search.

If the size of subtrees rooted at depth d is substantial, the cost of traversing a subtree is much higher than the overhead of jumping from one subtree to another. Assuming that subtrees at the same depth do not differ much in size, the cost of finding a goal node can be approximated by the number of subtrees one expects to explore. For search trees

with a single solution, the expected cost can be computed using the $P_{\text{goal}}(v)$ values of the nodes at the jump depth:

$$E_{\text{goal}}(d) = \frac{1}{2^d} \cdot \sum_{v=1}^{2^d} (\text{order}(v) \cdot P_{\text{goal}}(v)) \quad (4)$$

For a given depth d , subtrees are numbered chronologically from 1 to 2^d and $\text{order}(v)$ denotes the index at which subtree v will be visited in the specific search strategy. For example, the summation for ALDS using the goal node probabilities in Fig. 3 is:

$$\frac{1}{8} \cdot (1 \cdot .504 + 4 \cdot .056 + 3 \cdot .126 + 7 \cdot .014 + 2 \cdot .216 + 6 \cdot .024 + 5 \cdot .054 + 8 \cdot .006)$$

The table below shows the $E_{\text{goal}}(3)$ values based on the Fig. 3 probabilities:

	DFS	ILDS	DDS	ALDS
$E_{\text{goal}}(3)$	0.3375	0.31225	0.26375	0.26225

In case of multiple goal nodes, the 'expected' cost can be computed for a set of instances. Solve the instances using a search strategy and determine the average number of subtrees visited at depth d . Divide the average number by 2^d to obtain the average cost. We will denote this alternative by $E_{\text{goal}}^*(d)$.

4 Experiments

Two types of results are presented in this section: theoretical and experimental results. The theoretical results are based on a probabilistic model of heuristic tree search. Experiments have been performed with the look-ahead SAT solver *march* [12], the fastest solver on random 3-SAT benchmarks.

We compare several discrepancy-based search strategies on the theoretical model and on a dataset of random 3-SAT formulae. Additionally, two experiments were performed to determine how much ALDS could be improved.

4.1 Theoretical results

Based on the increasing heuristic probability observation [12] we created a model with just one goal node. In this model we assign the heuristic probability as follows (based on the observed $P_{\text{heur}}(v)$ values in [12]):

$$P_{\text{heur}}(v) = 0.56 + 0.015 \cdot \text{depth}(v) \quad (5)$$

Each leaf node v at depth 12 is assigned a goal node probability $P_{\text{goal}}(v)$ using the equations described in Section 3.1. This means P_{goal} is calculated by multiplying the heuristic probabilities of the left and right children leading to that leaf node, starting at the root with $P_{\text{heur}}(\text{root}) = 1$. So, similar to the tree in Fig. 3, only using a much larger tree.

In practice search trees contain multiple goal nodes, but no generality is lost by putting just one goal node in the search tree of our model. This will only cause expected cost to find a goal node and other numerical results to be a little larger, but this will not favour any search strategy in particular. The difference in numerical results is acceptable because this model is only used to compare the search strategies to each other.




Fig. 9. Theoretical fraction of instances unsolved versus number of leaf nodes explored on a model search tree of depth 12 containing a total of 4096 leaf nodes.

Because the goal node probability of each subtree is defined by our own model, the optimal order in which to search the subtrees is to go from high to low goal probability. The expected fraction of the tree that has to be searched before a goal node is reached is the area below each graph (see Fig. 9). ALDS performs best on the model. In addition, the difference between ALDS and the optimal search order is quite small.

4.2 Satisfiability results

For the experimental results we used the look-ahead SAT solver **march** [12] with the recursive weight heuristic w_3^x as direction heuristic (see Section 2.3). The dataset for the experiments consisted of 20146 satisfiable random 3-SAT instances with 350 variables and 1491 clauses. The clauses-to-variables ratio is 4.26, which is the ratio where the probability of generating a satisfiable instance is about 50%, known as the phase transition density.

For each instance, the complete search tree was explored to find out which of the 4096 subtrees at depth 12 contained solutions. A depth of 12 was chosen to keep the data

compact enough for practical use, but still perform discrepancy search on a significant part of the search tree. On average each instance contained 17.2 satisfiable subtrees.

Results from the experiments are shown in Fig. 10. The vertical axis shows the fraction of problems that has not been solved yet. The horizontal axis displays the number of subtrees that have been explored. Compared to the theoretical results, these lines decrease faster. This can be explained by the fact that there is an average of 17.2 goal nodes per instance in this experiment compared to a single goal node in the theoretical model. The $E_{\text{goal}}^*(12)$ values correspond to the area below the graphs.




Fig. 10. Fraction of instances unsolved versus explored search space, expressed by the number of subtrees explored on depth 12, for the 3-SAT instances.

4.3 Analysis

Similar to Section 4.1, we want to demonstrate that ALDS performs close to optimal on the dataset of random 3-SAT instances. Yet, due to multiple satisfying subtrees per instance, it is hard to determine the performance of the optimal search strategy. To approximate the optimal search strategy, we construct a *Greedy* search strategy. The Greedy search strategy is introduced in [12], and is constructed as follows:

- Select the subtree in which most instances from the dataset have at least one solution. This subtree is next to be visited in this specific Greedy search strategy.
- Remove from consideration all the instances in which the selected subtree has at least one solution.

- Repeat above steps until all instances are removed from consideration.
- The subtrees that have not been ordered yet, are placed at the end in ALDS order.

The construction of the Greedy search strategy requires a set of instances as input. We let $\text{Greedy}(S)$ denote the Greedy search strategy that has been constructed with input set S . Any Greedy search strategy will perform very well for the given input set. To determine whether or not Greedy could actually make a good *generalized* search strategy, the dataset has been split into two parts, which we will call part A and part B . When $\text{Greedy}(A)$ is applied to part A , the $E_{\text{goal}}^*(12)$ performance is very strong, as expected. However, when $\text{Greedy}(A)$ is applied to part B of the dataset, the result is worse than using the ALDS search method. This can be observed in Fig. 11.




Fig. 11. Analysis of Greedy strategies versus ALDS. Each of the pictures represents a distinct half of the dataset. The Greedy order used to create the top line in one picture is the order used to create the lower line in the other picture.

Furthermore, the $\text{Greedy}(S)$ lines on dataset S in Fig. 11 reach the horizontal axis after about 1024 of the 4096 subtrees. This is due to the limited size of the experiment. On a hypothetical collection of all possible random 3-SAT instances, this line could not reach the horizontal axis this quickly. To summarize: the Greedy search strategy benefits greatly from the limited size of the dataset. Yet, it seems unlikely that a greedy search strategy can be converted into a generalized search strategy in such a way that it will outperform ALDS.

In our second analysis of the performance of ALDS, we experimented with various linear models for the $P_{\text{heur}}(v)$ values, see (6). Our motivation to use this linear model is that we observed that the model (5) is an accurate approximation of the $P_{\text{heur}}(v)$ values using *march* with the w_3^* heuristic on random 3-SAT formulae.

$$P_{\text{heur}}(v) = y + x \cdot \text{depth}(v) \quad (6)$$

We used these linear models to construct search strategies:

- Compute for all 4096 subtrees at depth 12 the $P_{\text{goal}}(v)$ values
- Visit the subtrees based on the $P_{\text{goal}}(v)$ values from high to low

The objective is to determine which linear model results in the smallest $E_{\text{goal}}^*(12)$ value. Because we only use the $P_{\text{goal}}(v)$ values to determine the order of the subtrees, parameter y can be fixed. We set it to $y = 0.56$ as in (5). The table below shows the results:

x	0.0050	0.0055	0.0060	0.0065	0.0070	0.0075	0.0080
$E_{\text{goal}}^*(12)$	0.035331	0.035325	0.035327	0.035340	0.035341	0.035372	0.035390
x	0.0090	0.0100	0.0110	0.0120	0.0130	0.0140	0.0150
$E_{\text{goal}}^*(12)$	0.035434	0.035447	0.035561	0.035735	0.035840	0.035974	0.036057

Two interesting things can be concluded from these experiments. First, the optimal value for the x parameter (0.0055) is much smaller than the value that matches the observed data (0.0150). Second, the optimal $E_{\text{goal}}^*(12)$ value (0.03525) is hardly smaller than the cost of ALDS (0.036114). It will be difficult to construct a search strategy based on a linear model that will outperform ALDS, because the overhead required to implement a more complex strategy will probably be more costly than the small reduction in $E_{\text{goal}}^*(12)$ can compensate.

5 Conclusion

We introduced the recursive weight heuristic, a branching heuristic for look-ahead SAT solvers. The solver `march` equipped with this heuristic performs stronger than any other solver on unsatisfiable random 3-SAT formulae. It won the gold medal in this category at the SAT 2009 competition.

To capitalize on the recursive weight heuristic on satisfiable instances as well, we presented advanced limited discrepancy search (ALDS). Theoretical and practical results show that ADSL outperforms alternative complete search strategies on satisfiable random 3-SAT formulae. We show that on these instances, ALDS in combination with the recursive weight heuristic traverse the search tree in a near-optimal order.

References

1. Kullmann, O.: 7. [13] 205–244
2. Gomes, C.P., Selman, B., Crato, N., Kautz, H.: Heavy-tailed phenomena in satisfiability and constraint satisfaction problems. *J. Autom. Reason.* **24**(1-2) (2000) 67–100
3. Harvey, W.D., Ginsberg, M.L.: Limited discrepancy search, Morgan Kaufmann (1995) 607–613
4. Li, C.M., Anbulagan: Heuristics based on unit propagation for satisfiability problems. In: *IJCAI* (1). (1997) 366–371
5. Li, C.M.: A constraint-based approach to narrow search trees for satisfiability. *Information processing letters* **71**(2) (1999) 75–80
6. Dubois, O., Dequen, G.: A backbone-search heuristic for efficient solving of hard 3-SAT formulae. In Nebel, B., ed.: *IJCAI*, Morgan Kaufmann (2001) 248–253
7. Korf, R.E.: Improved limited discrepancy search. In: *Proceedings of AAAI-96*, MIT Press (1996) 286–291

8. Walsh, T.: Depth-bounded discrepancy search. In: IJCAI'97: Proceedings of the Fifteenth international joint conference on Artificial intelligence, San Francisco, CA, USA, Morgan Kaufmann Publishers Inc. (1997) 1388–1393
9. Davis, M., Logemann, G., Loveland, D.: A machine program for theorem-proving. *Commun. ACM* **5**(7) (1962) 394–397
10. Heule, M.J.H., van Maaren, H.: 5. [13] 155–184
11. Kullmann, O.: Investigating the behaviour of a SAT solver on random formulas. Technical Report CSR 23-2002, University of Wales Swansea, Computer Science Report Series (<http://www-compsci.swan.ac.uk/reports/2002.html>) (2002) 119 pages.
12. Heule, M.J.H., van Maaren, H.: Whose side are you on? finding solutions in a biased search-tree. Technical report, Proceedings of Guangzhou Symposium on Satisfiability In Logic-Based Modeling (2006)
13. Biere, A., Heule, M.J.H., van Maaren, H., Walsh, T., eds.: Handbook of Satisfiability. Volume 185 of Frontiers in Artificial Intelligence and Applications. IOS Press (2009)